

LISTING OF THE CLAIMS:

1. (Currently Amended) A method for finding optimal configurations of one or more clusters of resources given a set of constraints and policies, including manually/semi-automatically/automatically, and concurrently discovering resources, resource groups, leaving of resources availability and unavailability of the resources and resource groups, and determining their dynamic dependency dependencies and other configuration information such as capacities capacity and equivalency quality of service, cluster policies and changes thereto[,] in as applied to a network of resources, at cluster initialization and dynamically during cluster operation, along with for supporting a seamless startup and shutdown of the cluster and all its components clusters of resources according to the current policies and demands on the cluster resources and services.
2. (Original) A method to automatically build and incrementally manipulate an object database of resources from a discovered set of resources and their dependencies such that the database mirrors, digests, virtually centralises, partitions, and optimises a summary the cluster state, and ensures atomicity of changes, and enables linearisable views of the cluster with benefits.
3. (Original) A method for automatic invocation of one or more optimization solvers via a standard interface upon detecting state changes in the system, including an automatic representation of resource dependencies as a constraint graph and translating the constraint graph into a data structure suitable for any of a branch-and-bound type of search, a mathematical programming method, and an evolutionary optimisation method, upon receiving an external event.

4. (Original) A method for preprocessing events for optimisations and decisions including clubbing multiple enabled events, disjoint region or island by island processing, local optimisations, and creating the solution directly by bypassing the optimizer module.
5. (Original) A method for postprocessing a solution including its representation as a dependency graph with an implicit partial order for deployment of the solution, a translation of the dependency graph into multiple executable instructions and automatic and fault tolerant execution of such a plan of action, and a provision for auxiliary solutions in response to a failed execution of an instruction.
6. (Original) A method for partitioning a dependency graph with an implicit partial order into one or more disjoint regions for simultaneous, ordered, automatic, and fault-tolerant execution, with execution involving wiring up the repository with ordered commands, and with further support included for creation of artificial events for full island reevaluation in response to a poor deployment of a dependency graph.
7. (Original) A method for concurrent and pipelined handling of asynchronous state changes among resources in a distributed system, including a scheduler and atomic tasks with no permission for task preemption, parallel and pipelined execution of modules and their connecting channels, efficiency including no thread spinning, and a reduced number of state definitions, state transitions, and simple synchronisation such that system response is predictable.

8. (Original) A method for concurrent management of a distributed system and simulation of the same system using the same apparatus.

9. (Currently Amended) Apparatus for finding optimal configurations of one or more clusters of resources given a set of constraints and policies, including means for manually/semi-automatically/automatically, and concurrently discovering resources, resource groups, leaving-of resources availability and unavailability of the resources and resource groups, and means for determining their dynamic dependency dependencies and other configuration information such as capacities capacity and equivalency quality of service, cluster policies and changes thereto[,] in as applied to a network of resources, at cluster initialization and dynamically during cluster operation, along with for supporting a seamless startup and shutdown of the cluster and all its components clusters of resources according to the current policies and demands on the cluster resources and services.

10. (Original) Apparatus to automatically build and incrementally manipulate an object database of resources from a discovered set of resources and their dependencies such that the database mirrors, digests, (virtually) centralises, partitions, and optimises a summary the cluster state, and ensures atomicity of changes, and enables linearisable views of the cluster with benefits like a quick clubbing of multiple events and updates.

11. (Original) Apparatus for automatic invocation of one or more optimization solvers via a standard interface upon detecting state changes in the system, including an automatic representation of resource dependencies as a constraint graph and translating the constraint graph

into a data structure suitable for any of a branch-and-bound type of search, a mathematical programming method, and an evolutionary optimisation method, upon receiving an external event.

12. (Original) Apparatus for preprocessing events for optimisations and decisions including clubbing multiple enabled events, disjoint region or island by island processing, local optimisations, and creating the solution directly by bypassing the optimizer module.

13. (Original) Apparatus for postprocessing a solution including its representation as a dependency graph with an implicit partial order for deployment of the solution, a translation of the dependency graph into multiple executable instructions and automatic and fault tolerant execution of such a plan of action, and a provision for auxiliary solutions in response to a failed execution of an instruction.

14. (Original) Apparatus for partitioning a dependency graph with an implicit partial order into one or more disjoint regions for simultaneous, ordered, automatic, and fault-tolerant execution, with execution involving wiring up the repository with ordered commands, and with further support included for creation of artificial events for full island reevaluation in response to a poor deployment of a dependency graph.

15. (Original) Apparatus for concurrent and pipelined handling of asynchronous state changes among resources in a distributed system, including a scheduler and atomic tasks with no permission for task preemption, parallel and pipelined execution of modules and their connecting

channels, efficiency including no thread spinning, and a reduced number of state definitions, state transitions, and simple synchronisation such that system response is predictable.

16. (Original) Apparatus for concurrent management of a distributed system and simulation of the same system using the same apparatus.

17. (Currently Amended) A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for assisting a decision support system to find optimal configuration of one or more clusters of resources given a set of constraints and policies, said method steps comprising manually/semi-automatically/automatically, and concurrently discovering resources, resource groups, ~~leaving of resources~~ availability and unavailability of the resources and resource groups, and determining their dynamic dependency dependencies and other configuration information such as ~~capacities~~ capacity and equivalency quality of service, cluster policies and changes thereto[,] in as applied to a network of resources, at cluster initialization and dynamically during cluster operation, along with for supporting a seamless startup and shutdown of the cluster and all its components clusters of resources according to the current policies and demands on the cluster resources and services.

18. (Original) A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for assisting a decision support system, said method steps comprising automatically building and incrementally manipulating an object database of resources from a discovered set of resources and their dependencies such that

the database mirrors, digests, virtually centralises, partitions, and optimises a summary the cluster state, and ensures atomicity of changes, and enables linearisable views of the cluster with benefits.

19. (Original) A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for assisting a decision support system, said method steps comprising automatically invoking one or more optimization solvers via a standard interface upon detecting state changes in the system, including an automatic representation of resource dependencies as a constraint graph and translating the constraint graph into a data structure suitable for any of a branch-and-bound type of search, a mathematical programming method, and an evolutionary optimisation method, upon receiving an external event.

20. (Original) A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for preprocessing events for optimizations and decisions, said method steps including clubbing multiple enabled events, disjoint region or island by island processing, local optimizations, and creating the solution directly by bypassing the optimizer module.